

Abstract of the Disclosure

A system and method for improving the response time of text-to-speech synthesis utilizes “triphone contexts” (i.e., triplets comprising a central phoneme and its immediate context) as the basic unit, instead of performing phoneme-by-phoneme synthesis. The method comprises a method of generating a triphone preselection cost database for use in speech synthesis, the method comprising 1) selecting a triphone sequence $u_1 - u_2 - u_3$, 2) calculating a preselection cost for each 5-phoneme sequence $u_a - u_1 - u_2 - u_3 - u_b$, where u_2 is allowed to match any identically labeled phoneme in a database and the units u_a and u_b vary over the entire phoneme universe and 3) storing a group of the selected triphone sequences exhibiting the lowest costs in a triphone preselection cost database.